

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Systems and Methods for Filter Processing  
Using Hierarchical Data and Data Structures**

Inventors:

Geary L. Eppley

Umesh Madan

David Wortendyke

ATTORNEY'S DOCKET NO. MS1-1918US

## **TECHNICAL FIELD**

The systems and methods described herein generally relate to data processing using filters and, more particularly, to systems and methods for filter processing using hierarchical data associated with the filters and hierarchical data structures.

## **BACKGROUND**

Computing systems - i.e. devices capable of processing electronic data such as computers, telephones, Personal Digital Assistants (PDA), etc. - communicate with other computing systems by exchanging data messages according to a communications protocol that is recognizable by the systems. A system utilizes filters to analyze messages that are sent and/or received by the system and to determine if and how the messages will be processed further.

Filters are used to process a diverse assortment of documents and messages. Data processed by filters is frequently arranged in a hierarchical manner, generally as a result of an addressing or naming system that is used to uniquely identify or locate the data. For example, Internet resources are named according to Universal Resource Identifiers (URI) that reflect a hierarchical structure. The URI "http://www.mycorp.com/products" identifies a web page called "products" that is hosted on "www.mycorp.com" and is organized according to a particular URI scheme ("http"). In this particular example, "http" is a root of the hierarchy, "www.mycorp.com" is a subordinate of the root, and "/products" is a subordinate of "www.mycorp.com".

In addition to Internet resources, a distributed system frequently utilizes hierarchical namespaces as a natural way to organize and express the topology of

1 system resources such as servers, printers, etc. For example, the following  
2 hierarchical name scheme is a natural and convenient way of organizing MyCorp's  
3 web servers:

4 **`http://www.mycorp.com/machines/wa/seattle/server1.`**

5 In this particular example, the namespace utilizes the hierarchical structure  
6 to identify particular resources ("machines") in particular geographical areas  
7 ("wa" and "seattle"), which is a common scheme.

8 Whenever data is organized hierarchically, decision logic used to process  
9 the data also develops a hierarchical flavor. Thus, in the previous example, a  
10 corporation ("MyCorp") may develop a set of data processing rules that apply to  
11 machines regardless of the city and state in which the machines are located by  
12 orienting the rules about the name prefix "`http://www.mycorp.com/machines`".  
13 Another set of data processing rules that apply specifically to machines located in  
14 Washington would utilize the name prefix `http://www.mycorp.com/machines/wa`".

15 A common application that uses processing rules based on name prefixes is  
16 message routing. Message routers store route tables that are comprised of filters  
17 and destinations, a destination being associated with each filter. To route a  
18 message, a router matches the message against its internal filter table to produce a  
19 set of matching filters. The router then forwards the message to the destinations  
20 associated with each matching filter.

21 Such a router is relatively inefficient because the message router must try  
22 every input against each filter stored in the route table. Since route tables typically  
23 contain thousands of entries, the process of comparing an input with each and  
24 every filter can consume significant time and resources. If the number of  
25 comparisons is reduced, then the process would become more efficient.

1 Another problem is that, since multiple routes (i.e. filters) in a route table  
2 have overlapping scope, it is not uncommon for an input to match multiple filters.  
3 All matching filters that match a given input are returned in no particular order. If  
4 an order can be established and the results returned according to that order, then  
5 processing could be streamlined even further.

## 6 7 SUMMARY

8 At least one implementation described herein relates to filter processing  
9 according to hierarchical data associated with the filters. The filters are  
10 maintained according to a tree structure (i.e. a filter hierarchy) that is used to  
11 match inputs with filters. A filter hierarchy is an in-memory tree of string  
12 segments where each node of the tree corresponds to a string segment of the  
13 hierarchical data and references zero or more filters. Each string segment of a  
14 filter path corresponds with a level in the tree. When an input is received, a path  
15 (i.e. the hierarchical data) associated with the input is parsed into string segments  
16 and the tree is traversed according to the string segments to locate nodes that  
17 match the input (primary matching). The input is then compared to filters  
18 referenced by the matching nodes to locate filters that match the input (secondary  
19 matching). As a result, the input is tested against significantly fewer filters and  
20 greater efficiency is achieved over typical systems.

21 Depending on the particular implementation, the filter hierarchy may be  
22 traversed from top to bottom or bottom to top in the matching process. If matches  
23 deeper in the hierarchy take precedence over matches higher up, then matching  
24 begins at the bottom of the hierarchy and proceeds toward the top of the hierarchy.  
25 Conversely, if shallow matches are preferred, the matching begins at the top and

1 proceeds downward. No matter which direction the matching proceeds, matching  
2 filters are returned in an order in which they were encountered in the hierarchy.  
3 As a result, subsequent processing of the matching filters is more efficient.

4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

## **BRIEF DESCRIPTION OF THE DRAWINGS**

A more complete understanding of exemplary systems and methods described herein may be had by reference to the following detailed description when taken in conjunction with the accompanying drawings wherein:

Fig. 1 is an illustration depicting filter storage according to a hierarchical arrangement.

Fig. 2a is a diagram of an exemplary tree data structure representation.

Fig. 2b is a depiction of a message having multiple segments.

Fig. 3 is a block diagram of an exemplary computer system.

Fig. 4 is a flow diagram that depicts a methodological implementation for a primary matching process as described herein.

Fig. 5 is a flow diagram that depicts a methodological implementation for a secondary matching process as described herein.

Fig. 6 is a diagram of an exemplary computing environment in which the implementations described herein may operate.

## **DETAILED DESCRIPTION**

The present disclosure relates to organizing filters hierarchically and processing hierarchical data according thereto. Filters are stored in a hierarchical data structure that is tiered according to string segments contained in a segment path associated with the filters and/or the hierarchical data. The systems, methods and computer-readable media described herein provide a way in which a number of filters against which an input is tried is significantly reduced, thereby saving processing time and system resources.

Filters are stored in a tree structure that includes zero or more filters at each node. Each node represents a string segment in hierarchical data (such as a

namespace) and each level of the tree corresponds to a string segment position in the hierarchical data. When an input is received, hierarchical data associated with the input is parsed into string segments and the tree is traversed according to the string segments. When one or more nodes that correspond with the input are located, any filters associated with the nodes are compared with the input to determine if the input matches any of the filters.

By trying the input only against filters contained in a subset of the tree nodes instead of all the filters stored in the system, the number of tries is significantly reduced - on the order of hundreds or thousands. As such, the presently described systems and methods are an improvement over the prior art.

In addition, matching nodes are encountered in a particular order depending on a traversal algorithm that is utilized. For example, if a top-down traversal algorithm is used, nodes will be encountered according to a length of an input prefix, with a shortest prefix being encountered first and a longest prefix being encountered last. Matching nodes can be returned in the order in which they were encountered thus simplifying subsequent processing of the matching nodes.

In at least one implementation, a filter engine uses XPath filters and input data is structured according to the eXtensible Markup Language (XML) data format. This combination works particularly well since these formats are hierarchical in nature. However, the techniques described herein may also be utilized with other hierarchical formats.

### **Exemplary Data Structure**

**Fig. 1** is an illustration depicting filter storage according to a hierarchical arrangement. Fig. 1 shows an exemplary filter hierarchy depicted by a filter tree 100. The filter tree 100 represents filters associated with two pathnames and

1 several subsets thereof. A first pathname represented by the filter tree 100 is  
2 soap://xyz.com:35. A second pathname represented by the filter tree 100 is  
3 soap://mycorp.com/a/b/c. One subset of the second pathname is  
4 soap://mycorp.com/a. Other subsets are included in the second pathname.

5 A root node 102 includes a first segment of each pathname, namely,  
6 "soap://". According to Fig. 1, all inputs received for the filter tree 100 will begin  
7 with this first segment. It is noted that a root node may be a blank node or a null  
8 node that represents a unity of dissimilar first segments of inputs that may be  
9 received for a particular filter tree. For example, if the filter tree 100 were also  
10 configured to received inputs having a first segment of "http://" then a blank root  
11 node would exist and that root node would include a first child node - "http://" -  
12 and a second child node - "soap://". However, that is not the case depicted in the  
13 present example.

14 In the present example, the root node 102 includes a first child node 104  
15 and a second child node 106. Each child node 104, 106 represents a unique  
16 segment that appears in a similar position of an input. Here, the first child node  
17 104 references a segment from the first pathname, namely "xyz.com". An input  
18 that includes the string "soap://xyz.com" would match the first child node 104.  
19 The first child node 104 also has a child node 105 that references a second  
20 segment from the first pathname, ":35" (a port identifier). An input that includes  
21 the string "soap://xyz.com" would match child node 105.

22 Similarly, the second child node 106 references a segment from the second  
23 path name ("mycorp.com") that is located in a similar position within the  
24 pathname as "xyz.com". Since both segments are located in a similar position  
25 within the pathname, the segments are represented in a same level in the filter tree



1 100. An input that includes the string "soap://mycorp.com" would match the  
2 second child node 106.

3 The second string ("soap://mycorp.com/a/b/c") has five string segments  
4 following the segment represented by the second child node 106  
5 ("soap://mycorp.com"). The five string segments are ":80" (the default port for  
6 HyperText Transmission Protocol (http) which is used by soap (simple object  
7 access protocol; since it is a default port, it does not appear in the string))  
8 messages, "/", "a", "/b" and "/c". Each of these string segments is represented by  
9 a node in the filter tree 100. Node 108 represents ":80" and node 110 represents  
10 "/". Node 110 is subordinate to node 108 since ":80" appears before "/" in the  
11 associated string.

12 Since "a" appears before "/b" and "/c", a node 112 representing the  
13 segment "a" is located subordinate to the node 110 representing the "/" string  
14 segment. The string segment "/b" appears subsequent to the string segment "a" so  
15 a node 114 corresponding to the "/b" string segment is subordinate to the node 112  
16 that represents the "a" string segment. Similarly, a node 116 representing the "/c"  
17 string segment is shown subordinate to the node 114 representing the "/b" string  
18 segment since the "/c" string segment occurs subsequent to the "/b" string  
19 segment.

20 Filters are stored and referenced by individual nodes in the filter tree 100.  
21 A particular node in the filter tree 100 may have multiple filters associated with it,  
22 but any particular node may not have any filters associated therewith. In addition,  
23 a node may have application data associated therewith that is a set of instructions  
24 that are executed if a filter associated with the node is satisfied by an input.  
25

1 In the present example, node 114 has a filter 118 associated with it that  
2 includes a condition of "Priority < 4" and application data that is a server  
3 designation of "person.com". When an input including "soap://mycorp.com/a/b"  
4 is tried against the filter tree 100, it will match node 114. The application data of  
5 filter 118 will then be executed if the input satisfies the condition (i.e. if the input  
6 has a priority that is less than four) and the input will be forwarded to  
7 "person.com".

8 Similarly, node 116 has a filter 120 associated therewith. Filter 120  
9 includes a condition of "Priority = 1" and application data that is a server  
10 designation of "group.com". When an input including "soap://mycorp.com/a/b/c)"  
11 is tried against the filter tree 100, it will match node 116. The application data of  
12 filter 120 will then be executed if the condition - a priority that is equal to one - is  
13 satisfied. The input will be forwarded to the "group.com" server.

14 To better illustrate the prefix matching nature of the described  
15 implementations, consider the following input:

16 **soap://mycorp.com/a/b/d/e Priority = 3**

17 Here, the tree would be traversed down to node 114 ("b") and the filter 118  
18 associated with node 114 would be tested against the input since no further  
19 segment of the input matched a node lower than "b". Since the priority of the  
20 input is less than four (i.e. the filter evaluates to true), the input would be  
21 forwarded to person.com. However, in some implementations, it may be  
22 advantageous to consider only exact matches. In this instance, there is no exact  
23 match.

24 Node 104 also has a filter 105 associated with it. The application data  
25 associated with filter 105 only includes a server designation ("corp.com") but no

1 conditional rule. This situation may be viewed as filter 105 being the server  
2 designation, or that filter 105 has no filter but has an instruction set that is  
3 executed when an input matches node 104. Furthermore, for representational  
4 purposes, it may be said that filter 105 includes a null filter with application data  
5 of a server designation. An input including “soap://xyz.com” that is tried against  
6 the filter tree will simply be forwarded to “corp.com”.

### 7 **Exemplary String Segment and Associated Data Structure**

8 **Fig. 2a and Fig. 2b** illustrate how filter tree levels correspond to data item  
9 segments. Fig. 2a is a diagram of an exemplary hierarchical data structure  
10 representing a filter tree 200. Fig. 2b is a depiction of a data item 202 having  
11 segments ordered according to a hierarchical scheme. The filter tree 200 includes  
12 multiple levels that each correspond to a segment of the data item 202.

13 The filter tree 200 includes a first level 210a that has a root node; a second  
14 level 212a that contains child nodes of the root node; a third level 214a that is  
15 subordinate to the second level 212a; a fourth level 216a that is subordinate to the  
16 third level 214a; and a fifth level 218a that has bottom nodes and is subordinate to  
the fourth level 216a.

17 The data item 202 includes a first segment 210b (“soap://”) that  
18 corresponds with the root node in the first level 210a of the filter tree 200. The  
19 data item 202 also includes a second segment 212b (“mycorp.com”) that occurs  
20 subsequent to the first segment 210a and that corresponds with the second level  
21 212a of the filter tree 200. In other words, the second segment 212b would be  
22 associated with a node in the second level 212b of the filter tree 200.

23 The data item 202 also includes a third segment 214b (“:80”), a fourth  
24 segment 216b (“/”) and a fifth segment 218b (“a”) that correspond with the third  
25 level 214a, fourth level 216a and fifth level 218a, respectively, of the filter tree

1 200. In this and other examples, the order in which a segment appears in a data  
2 item determines the level in a filter tree that is associated with the segment.

3 The advantages of such a corresponding relationship between hierarchical  
4 data items and hierarchical filter data structures will be seen below, when further  
5 aspects of the present systems and methods are described in greater detail.

### 6 Exemplary Computer System

7 **Fig. 3** is a simplified block diagram of a computer system 300 showing key  
8 components described herein. It is noted that the computer system 300 includes  
9 other components than those shown, i.e. components typically found in computer  
10 systems. However, only elements particularly related to the systems and methods  
11 described herein are depicted. It is also noted that the elements shown in Fig. 3  
12 may be implemented as hardware components, software modules or a combination  
thereof.

13 The system 300 includes memory 302 that stores a filter tree 304  
14 representation that has a tree logic component 306 and a filter tables component  
15 308. The tree logic component 306 includes the logic required to build and  
16 maintain the filter tree 304 and the filter tables component 308 includes one or  
17 more filters referenced by the filter tree 304.

18 It is noted that although the filter tables 308 are shown resident in memory  
19 302, filter tables used by the system could be stored in remote memory and  
20 swapped into memory 302 as needed. It is also noted that it is possible that a  
21 message being streamed into the system 300 may never be completely resident in  
22 the memory 302 at one time. In such a case, filtering could be performed on a  
23 front-end of the message (such as a header section) and the stream redirected to  
another process based on the result of the filtering.

24 The memory 302 also stores a message input module 310 and a message  
25 processing module 312. The memory further includes two matching modules,

1 namely a primary matching module 314, a secondary matching module 316 and a  
2 nodeset 318.

3 The message input module 310 is configured to receive a data transmission  
4 containing an input to try against filters in the filter tree 304. The message  
5 processing module 312 is configured, inter alia, to parse the input into distinct  
6 portions that correspond to a hierarchy of the filter tree 304. In the case of a  
7 message, the message processing module 312 is configured to parse the message  
8 into its constituent segments for subsequent matching procedures.

9 Other functions attributable to the message processing module 312 (such as  
10 precedence processing, etc.) will be discussed in greater detail below, with respect  
11 to subsequent figures.

12 The primary matching module 314 is configured to locate nodes in the filter  
13 tree 304 that match segments contained in an input message. The secondary  
14 matching module 316 is configured to compare the input message used in the  
15 primary matching procedure to one or more filters referenced by the matching  
16 nodes identified in the primary matching procedure.

17 By separating the matching process into primary and secondary procedures  
18 performed by the primary matching module 314 and the secondary matching  
19 module 316, respectively, a particular input must only be tried against filters  
20 identified by the secondary matching module to determine all matches between the  
21 input and the filters in the filter tree. By eliminating a vast majority of matching  
22 attempts between the input and the non-matching filters, the efficiency of the  
23 matching process is greatly enhanced.

#### 24 **Methodological Implementation: Exemplary Primary Matching**

25 **Fig. 4** is a flow diagram 400 depicting an exemplary primary matching  
process to locate one or more nodes in a filter tree that correspond to segments

1 included in an input having a hierarchical structure. In the following discussion,  
2 continuing reference will be made to the elements and the corresponding reference  
3 numerals included in previous figures.

4 The following example relates to a messaging system that receives input  
5 messages and compares the input to several filters to determine filters that match  
6 the input. However, it is noted that other types of systems can utilize the  
7 methodology presented herein.

8 Also, it is noted that function may be attributed to a particular element for  
9 purposes of this discussion. However, such attribution does not necessarily limit  
10 the function to that particular element. Any function discussed in the following  
11 methodological implementation may be performed by any component capable of  
12 performing such a function.

13 At block 402, an input is received by the message input module 310 (Fig.  
14 3). The message input module 310 identifies a hierarchical string contained in the  
15 input (block 404) and passes the string to the message processing module 312,  
16 which parses the string into identifiable string segments at block 406.

17 At block 408, the primary matching module 314 traverses the filter tree and  
18 compares the segments (see Fig. 2a) with filter tree nodes in corresponding filter  
19 tree levels (see Fig. 2b) to determine if any filter tree nodes match the hierarchical  
20 string of the input message. The details of the filter tree traversal are not described  
21 in detail herein. Tree traversal algorithms are known in the art and may be used  
22 with the present systems and methods.

23 The traversal begins with a root node and progresses through each filter tree  
24 branch as long as matching nodes are encountered. Once a non-matching node is  
25 encountered in a filter tree branch, further traversal of that branch is not required,

1 since no nodes that occur after the non-matching node can match the input string.  
2 This is one aspect of the present invention that increases efficiency of a filter  
3 engine process.

4 Additionally, new levels of the tree are not traversed if there is no string  
5 segment that corresponds to the new levels. For example, if a string only has three  
6 string segments, there is no need to traverse tree levels from a fourth level and  
7 below.

8 Any nodes discovered during the traversal that match a corresponding  
9 segment of the input string are added to the nodeset 318 (block 410). Further  
10 processing takes place using only the nodes included in the nodeset 318. This is  
11 another aspect of the present invention that preserves system resources, since filter  
12 matching is only performed on a subset of tree nodes.

13 In at least one implementation, only nodes that are associated with a filter  
14 (or more than one filter) are added to the nodeset. After a node is found to match  
15 an input, a check is made as to whether a filter is referenced by the node. If so,  
16 then the node is added to the nodeset; otherwise, it is not.

### 17 **Methodological Implementation: Exemplary Secondary Matching**

18 **Fig. 5** is a flow diagram 500 that depicts an exemplary methodological  
19 implementation of a secondary matching process. In the following discussion,  
20 reference will be made to elements and reference numerals shown in previous  
21 figures.

22 At block 502, the secondary matching module 316 identifies a first node  
23 stored in the nodeset 318. The secondary matching module 316 then identifies one  
24 or more filters associated with the node at block 504. For example, referring to  
25 the example shown in Fig. 1, if the first node in the nodeset 318 identifies node

1 114, then the secondary matching module 316 would identify filter 118 at block  
2 504.

3 At block 506, the secondary matching module 316 compares the input to  
4 the filter(s) identified in block 504 to determine if the input satisfies any of the  
5 filters. If a filter is found that satisfies the input ("Yes" branch, block 508), then  
6 the secondary matching module 316 notes the match at block 510, typically by  
7 caching a reference to the node, filter or action.

8 If there is no match ("No" branch, block 508) or after a match has been  
9 noted (block 510) there are more nodes in the nodeset ("Yes" branch, block 512),  
10 then a next node in the nodeset 318 is identified at block 514 and the process  
11 reverts to block 504 to identify node filters and compare the input to the filters.  
12 When there are no unprocessed nodes remaining in the nodeset 318 ("No" branch,  
13 block 512), the secondary matching module selects one or more actions from the  
14 matching nodes and executes the action(s) (block 516). The action(s) taken may  
15 include selecting only one action at random or selecting one or more of the actions  
16 based on a predetermined methodology.

17 The steps described above are but one part of a particular implementation  
18 of the secondary matching process wherein the actions of matching filters are  
19 executed as they are identified. In another implementation which may better suit  
20 some applications, the secondary matching module 316 executes an action  
21 associated with a filter upon determining that an input satisfies the filter. In other  
22 words, the actions are taken during the iteration through the nodeset instead of  
23 after all the nodes have been processed.

24 The above implementation might be used by an application that enforces  
25 security policy that is applied in a hierarchical manner. Each filter may correspond



1 to a policy that has to be enforced. Therefore, all matching filters would need to  
2 be identified.

3 In yet another implementation of the secondary matching process, the  
4 nodes in the nodeset 318 may be processed only until a match is identified. The  
5 processing may begin from the top or the bottom of the nodeset, i.e. the top or  
6 bottom of the filter tree. Such an implementation might be used starting at the  
7 bottom and working backwards for a routing system that defines default routes at  
8 higher nodes to catch messages that do not match more specific routes defined at  
9 lower levels.

10 Implementing the filter matching process as described above obviates the  
11 need to match an input against each and every filter in a filtering system. The only  
12 comparing of an input to individual filters occurs in the secondary matching  
13 process after the primary matching process has eliminated a significant number of  
14 filters from the possibility of matching the input. This saves significant system  
15 resources over other methods.

### 16 **Precedence Considerations**

17 One issue that arises with hierarchical data items such as namespaces is that  
18 one input may match more than one namespace or filter. For example, consider an  
19 enterprise that defines rules that apply to enterprise machines according to name  
20 prefixes (as in the example previously discussed).

21 Precedence rules would have to be defined by the enterprise to settle  
22 particular issues. For example, would rules for `http://www.mycorp.com/machines`  
23 always take precedence over `http://www.mycorp.com/machines/wa/seattle`?  
24 Would a rule defined for the prefix `http://www.mycorp.com/machines` apply to  
25 every machine whose name begins with that prefix? Does a rule specific to

1 http://www.mycorp.com/machines/wa/seattle take precedence over a more general  
2 rule for http://www.mycorp.com/machines/wa?

3 Typically, a precedence algorithm is based on either a longest prefix or a  
4 shortest prefix basis. A longest prefix precedence would mean that a more specific  
5 rule would take precedence over a more general rule. Accordingly, a shortest  
6 prefix precedence would give precedence to the more general rule.

7 As previously discussed, a common usage of name prefix based grouping  
8 of processing rules is in message routing systems. Message systems contain route  
9 tables made up of filters, with each filter having an associated destination. To  
10 route a message, a router matches the message against its internal filter table to  
11 produce a set of matching filters. It then forwards the message to the destinations  
12 associated with each matching filter.

13 Message routers route messages to Web services and locations, which are  
14 addressed and located using URIs. It is possible for multiple routes (filters) in a  
15 single message routing system to have overlapping scope. It is not uncommon for  
16 a message to match multiple routes.

17 **Example:** Referring back to Fig. 1, it can be seen that filter 116 represents  
18 the following rule (hereinafter "Route 1"): if a message "To:" begins with  
19 soap://mycorp.com/a/b and has a priority that is less than four (4), forward the  
20 message to person.com. Furthermore, filter 118 represents the following rule  
21 ("Route 2"): if a message "To:" begins with soap://mycorp.com/a/b/c and has a  
22 priority equal to one (1), route the message to group.com. Finally, filter 114  
23 represents a rule ("Route 3") of: if a message "To:" begins with soap://xyz.com,  
24 route the message to corp.com regardless of the priority.

1        Given an input message "To: soap://mycorp.com/a/b/c" having a priority of  
2 one (1), a router would select two routes: Route 1 (filter 116) and Route 2 (filter  
3 118). That in itself may be acceptable, but what if the router were required to pick  
4 a single route?

5        If a precedence rule mandated that a route with the shortest prefix be  
6 selected, then Route 1 would be the only route deemed to match the input  
7 message. This is because Route 1 has a shorter prefix than Route 2 -  
8 "soap://mycorp.com/a/b" vs. "soap://mycorp.com/a/b/c".

9        According to the methodologies present herein, a primary matching step  
10 could traverse the filter tree 100 to determine a set of nodes that (a) have filters  
11 and/or application data associated therewith; and (b) match the input. This step  
12 would derive node 110 and node 112 as matching nodes. Node 106 and node 108  
13 would not match because they do not have any filter or data associated with them.  
14 Node 104 would not match because the input message string segment does not  
15 match "soap://xyz.com".

16        A router or application would then apply a precedence rule to determine  
17 which specific filters should be tried (as part of the secondary matching process  
18 described above). Using the shortest prefix precedence rule, Route 1 (filter 116)  
19 would be selected because it has the shortest matching prefix, namely,  
20 "soap://mycorp.com/a/b" (node 110).

21        Although not clearly obvious in this particular example due to the limited  
22 number of filters shown, when a filter tree containing hundreds or thousands of  
23 filters is considered, the reduction in resource requirements become apparent.  
24 This technique is more efficient than a typical technique that would try the input  
25 against all filters in the filtering system.

## Exemplary Matching Algorithm with Precedence Considerations

The following exemplary matching algorithm is described in pseudo code to better demonstrate the techniques described herein. The following examples include algorithms for shortest prefix precedence and for longest prefix precedence.

```
internal HierarchyNodeCollection Match (string[ ] segments)
{
    HierarchyNodeCollection matches = new
    FilterHierarchyNodeCollection( );
    HierarchyNode node = this.Root;

    if (node.HasFilters)
    {
        matches.Add(node);
    }

    // Loop over each segment
    foreach (string segment in segments)
    {
        //For each segment, find a child node that has a matching
        segment
        HierarchyNode match = node.Find(segment);
        If (null != match && match.HasFilters)
        {
            matches.Add(match)
        }
        else
        {
            break; //no match. Traversal stops
        }
    }

    return matches.
}
```

The preceding pseudo code traverses a filter tree starting with the root node. The matching nodes are returned in the order they occur in the tree, starting at the

1 root and down. To give precedence to shortest prefix matches, the application  
2 must perform the following steps:

```
3
4 FilterCollection filters = null;
5 int i = 0;
6 do while (filters == null && i < matches.Length)
7 {
8     filters = matches[i].Match(input);
9     i = i + 1;
10 }
```

11 This selects the first matching node (the node having a shortest prefix) and  
12 asks the filter hierarchy to perform a second pass of the input data. The hierarchy  
13 will test all filters references by the first matching node against the input and  
14 return filters that match.

15 To give precedence to the longest prefix matches, the application performs:

```
16
17 FilterCollection filters = null;
18 int i = matches.Length-1;
19 do while (filters == null && i >= 0)
20 {
21     filters = matches[i].Match(input);
22     i = i - 1;
23 }
```

24 This selects the last matching node (the node having the longest prefix) and  
25 asks the filter hierarchy to perform a second pass on the input data.

To perform other precedence rules, an application can loop over the  
returned matches and perform whatever action is preferred.

### Adding and Removing Filters To/From a Filter Hierarchy

Adding filters to a filter hierarchy and removing filters therefrom follows  
similar logic as outlined above. To add a filter to a filter hierarchy, a hierarchically  
structured segment string that is associated with the filter is used to locate a  
nearest matching node in the filter hierarchy, i.e. a filter tree.

1       The filter tree is traversed according to the segment string, as explained  
2 above. If an existing node is found that matches the segment string, the associated  
3 filter is added to the node. In other words, the node is made to reference the new  
4 filter along with any other filters that the node may already reference.

5       If there is no existing node that matches the segment string exactly, the  
6 lowest level node in the filter tree that matches a segment in the segment string is  
7 located by traversing the filter tree. A child node is created for that lowest level  
8 node that corresponds with a string segment that is immediately subsequent to the  
9 last matching string segment.

10       If the child node (i.e. the new node) corresponds with a final string  
11 segment, then a filter and application data associated with the segment string is  
12 associated with the child node. If there is another subsequent string segment, then  
13 a grandchild node from the lowest level matching node is made subordinate to the  
14 new child node created above. If the grandchild node corresponds to a final string  
15 segment, then the grandchild node is made to reference the filter and application  
16 data associated with the segment string.

17       This process continues until a final string segment in the segment string has  
18 been accounted for.

19       The following pseudo code addresses the previous steps.

```
20 //  
21 // PSEUDO CODE  
22 // Traverses the tree, making sure that the tree has nodes  
23 // corresponding to the segments in segments[]  
24 // Returns the LAST - LEAF node in the tree - matching the last segment  
25 // This is the node with which a filter is associated  
  
public FilterHierarchyNode Add(string[] segments)  
{  
    // start traversing at the root of the tree  
    FilterHierarchyNode node = this.root;  
    FilterHierarchyNode newNode;  
  
    for (int i = 0; i < segments.Length; ++i)
```

```

1      {
2          string currentSegment = segments[i];
3
4          // Does the current node have a child whose segment matches
5          // the current one?
6          if (null == (nextNode = node.GetChild(currentSegment)))
7          {
8              // No. Create a new node
9              newNode = this.CreateNode(currentSegment);
10             node.AddChild(newNode);
11             // Keep traversing from the new node
12             node = newNode;
13         }
14         else
15         {
16             // A suitable node already exists. Continue traversal
17             node = nextNode;
18         }
19     }
20
21     return node;    // corresponds to segments[segments.Length - 1]
22 }

```

Removing a filter from a filter tree is accomplished by locating a node associated with the filter and removing any reference to the filter and/or application data from the node.

If the node associated with the filter does not reference any other filter or application data, and if the node does not have any subordinate nodes, then the node is removed from the filter tree. A parent node of the removed node may also be removed if the previous node removal removes a last child node from the parent node and if the parent node does not reference any filter or application data. This removal process can be repeated up the filter tree until a node is reached that has a child node or references a filter and/or application data.

### **Exemplary Computer Environment**

The various components and functionality described herein are implemented with a computing system. Fig. 6 shows components of typical example of such a computing system, i.e. a computer, referred by to reference

1 numeral 600. The components shown in Fig. 6 are only examples, and are not  
2 intended to suggest any limitation as to the scope of the functionality of the  
3 invention. Furthermore, the invention is not necessarily dependent on the features  
4 shown in Fig. 6.

5 Generally, various different general purpose or special purpose computing  
6 system configurations can be used. Examples of well known computing systems,  
7 environments, and/or configurations that may be suitable for use with the  
8 invention include, but are not limited to, personal computers, server computers,  
9 hand-held or laptop devices, multiprocessor systems, microprocessor-based  
10 systems, set top boxes, programmable consumer electronics, network PCs,  
11 minicomputers, mainframe computers, distributed computing environments that  
12 include any of the above systems or devices, and the like.

13 The functionality of the computers is embodied in many cases by computer-  
14 executable instructions, such as program modules, that are executed by the  
15 computers. Generally, program modules include routines, programs, objects,  
16 components, data structures, etc. that perform particular tasks or implement  
17 particular abstract data types. Tasks might also be performed by remote  
18 processing devices that are linked through a communications network. In a  
19 distributed computing environment, program modules may be located in both local  
20 and remote computer storage media.

21 The instructions and/or program modules are stored at different times in the  
22 various computer-readable media that are either part of the computer or that can be  
23 read by the computer. Programs are typically distributed, for example, on floppy  
24 disks, CD-ROMs, DVD, or some form of communication media such as a  
25 modulated signal. From there, they are installed or loaded into the secondary



1 memory of a computer. At execution, they are loaded at least partially into the  
2 computer's primary electronic memory. The invention described herein includes  
3 these and other various types of computer-readable media when such media  
4 contain instructions programs, and/or modules for implementing the steps  
5 described below in conjunction with a microprocessor or other data processors.  
6 The invention also includes the computer itself when programmed according to  
7 the methods and techniques described below.

8 For purposes of illustration, programs and other executable program  
9 components such as the operating system are illustrated herein as discrete blocks,  
10 although it is recognized that such programs and components reside at various  
11 times in different storage components of the computer, and are executed by the  
12 data processor(s) of the computer.

13 With reference to Fig. 6, the components of computer 600 may include, but  
14 are not limited to, a processing unit 602, a system memory 604, and a system bus  
15 606 that couples various system components including the system memory to the  
16 processing unit 602. The system bus 606 may be any of several types of bus  
17 structures including a memory bus or memory controller, a peripheral bus, and a  
18 local bus using any of a variety of bus architectures. By way of example, and not  
19 limitation, such architectures include Industry Standard Architecture (ISA) bus,  
20 Micro Channel Architecture (MCA) bus, Enhanced ISA (EISAA) bus, Video  
21 Electronics Standards Association (VESA) local bus, and Peripheral Component  
22 Interconnect (PCI) bus also known as the Mezzanine bus.

23 Computer 600 typically includes a variety of computer-readable media.  
24 Computer-readable media can be any available media that can be accessed by  
25 computer 600 and includes both volatile and nonvolatile media, removable and

1 non-removable media. By way of example, and not limitation, computer-readable  
2 media may comprise computer storage media and communication media.  
3 “Computer storage media” includes volatile and nonvolatile, removable and non-  
4 removable media implemented in any method or technology for storage of  
5 information such as computer-readable instructions, data structures, program  
6 modules, or other data. Computer storage media includes, but is not limited to,  
7 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,  
8 digital versatile disks (DVD) or other optical disk storage, magnetic cassettes,  
9 magnetic tape, magnetic disk storage or other magnetic storage devices, or any  
10 other medium which can be used to store the desired information and which can be  
11 accessed by computer 600. Communication media typically embodies computer-  
12 readable instructions, data structures, program modules or other data in a  
13 modulated data signal such as a carrier wave or other transport mechanism and  
14 includes any information delivery media. The term “modulated data signal”  
15 means a signal that has one or more of its characteristics set or changed in such a  
16 manner as to encode information in the signal. By way of example, and not  
17 limitation, communication media includes wired media such as a wired network or  
18 direct-wired connection and wireless media such as acoustic, RF, infrared and  
19 other wireless media. Combinations of any of the above should also be included  
20 within the scope of computer readable media.

21 The system memory 604 includes computer storage media in the form of  
22 volatile and/or nonvolatile memory such as read only memory (ROM) 608 and  
23 random access memory (RAM) 610. A basic input/output system 612 (BIOS),  
24 containing the basic routines that help to transfer information between elements  
25 within computer 600, such as during start-up, is typically stored in ROM 608.

1 RAM 610 typically contains data and/or program modules that are immediately  
2 accessible to and/or presently being operated on by processing unit 602. By way  
3 of example, and not limitation, Fig. 6 illustrates operating system 614, application  
4 programs 616, other program modules 618, and program data 620.

5 The computer 600 may also include other removable/non-removable,  
6 volatile/nonvolatile computer storage media. By way of example only, Fig. 6  
7 illustrates a hard disk drive 622 that reads from or writes to non-removable,  
8 nonvolatile magnetic media, a magnetic disk drive 624 that reads from or writes to  
9 a removable, nonvolatile magnetic disk 626, and an optical disk drive 628 that  
10 reads from or writes to a removable, nonvolatile optical disk 630 such as a CD  
11 ROM or other optical media. Other removable/non-removable,  
12 volatile/nonvolatile computer storage media that can be used in the exemplary  
13 operating environment include, but are not limited to, magnetic tape cassettes,  
14 flash memory cards, digital versatile disks, digital video tape, solid state RAM,  
15 solid state ROM, and the like. The hard disk drive 622 is typically connected to  
16 the system bus 606 through a non-removable memory interface such as data media  
17 interface 632, and magnetic disk drive 624 and optical disk drive 628 are typically  
18 connected to the system bus 606 by a removable memory interface such as  
19 interface 634.

20 The drives and their associated computer storage media discussed above  
21 and illustrated in Fig. 6 provide storage of computer-readable instructions, data  
22 structures, program modules, and other data for computer 600. In Fig. 6, for  
23 example, hard disk drive 622 is illustrated as storing operating system 615,  
24 application programs 617, other program modules 619, and program data 621.  
25 Note that these components can either be the same as or different from operating

1 system 614, application programs 616, other program modules 618, and program  
2 data 620. Operating system 615, application programs 617, other program  
3 modules 619, and program data 621 are given different numbers here to illustrate  
4 that, at a minimum, they are different copies. A user may enter commands and  
5 information into the computer 600 through input devices such as a keyboard 636  
6 and pointing device 638, commonly referred to as a mouse, trackball, or touch  
7 pad. Other input devices (not shown) may include a microphone, joystick, game  
8 pad, satellite dish, scanner, or the like. These and other input devices are often  
9 connected to the processing unit 602 through an input/output (I/O) interface 640  
10 that is coupled to the system bus, but may be connected by other interface and bus  
11 structures, such as a parallel port, game port, or a universal serial bus (USB). A  
12 monitor 642 or other type of display device is also connected to the system bus  
13 606 via an interface, such as a video adapter 644. In addition to the monitor 642,  
14 computers may also include other peripheral output devices 646 (e.g., speakers)  
15 and one or more printers 648, which may be connected through the I/O interface  
16 640.

17 The computer may operate in a networked environment using logical  
18 connections to one or more remote computers, such as a remote computing device  
19 650. The remote computing device 650 may be a personal computer, a server, a  
20 router, a network PC, a peer device or other common network node, and typically  
21 includes many or all of the elements described above relative to computer 600.  
22 The logical connections depicted in Fig. 6 include a local area network (LAN) 652  
23 and a wide area network (WAN) 654. Although the WAN 654 shown in Fig. 6 is  
24 the Internet, the WAN 654 may also include other networks. Such networking  
25

1 environments are commonplace in offices, enterprise-wide computer networks,  
2 intranets, and the like.

3 When used in a LAN networking environment, the computer 600 is  
4 connected to the LAN 652 through a network interface or adapter 656. When used  
5 in a WAN networking environment, the computer 600 typically includes a modem  
6 658 or other means for establishing communications over the Internet 654. The  
7 modem 658, which may be internal or external, may be connected to the system  
8 bus 606 via the I/O interface 640, or other appropriate mechanism. In a networked  
9 environment, program modules depicted relative to the computer 600, or portions  
10 thereof, may be stored in the remote computing device 650. By way of example,  
11 and not limitation, Fig. 6 illustrates remote application programs 660 as residing  
12 on remote computing device 650. It will be appreciated that the network  
13 connections shown are exemplary and other means of establishing a  
14 communications link between the computers may be used.

### 15 Conclusion

16 The systems and methods as described thus provide an efficient way to  
17 match inputs against multiple filters by reducing a number of filters that must be  
18 tried against the input. Furthermore, the techniques describe provide an efficient  
19 way to sort and group filters so that a precedence algorithm may be applied.

20 Although details of specific implementations and embodiments are  
21 described above, such details are intended to satisfy statutory disclosure  
22 obligations rather than to limit the scope of the following claims. Thus, the  
23 invention as defined by the claims is not limited to the specific features described  
24 above. Rather, the invention is claimed in any of its forms or modifications that  
25

1 fall within the proper scope of the appended claims, appropriately interpreted in  
2 accordance with the doctrine of equivalents.

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25